

WHAT IS CLAIMED IS:

1. A node for use in a multi-node system, the node comprising:

5 a plurality of active devices;

an interface to an inter-node network coupling nodes in the multi-node system;
and

10 an address network configured to convey address packets between the interface
and the plurality of active devices;

wherein an active device of the plurality of active devices is configured to send an
address packet on the address network to initiate a transaction to gain an
15 access right to a coherency unit;

wherein in response to the address packet, the interface is configured to send data
corresponding to the coherency unit to the active device if no other active
device in the node has an ownership responsibility for the coherency unit
20 and the coherency unit is in a modified global access state in the node.

2. The node of claim 1, wherein the interface includes storage for a plurality of
records, wherein each of the plurality of records corresponds to a respective one of a
plurality of address packets conveyed on the address network;

25

wherein the interface is configured to search the storage for records specifying the
coherency unit in response to the interface receiving a request for access to
the coherency unit from another node via the inter-node network;

wherein the interface is configured to send the data corresponding to the coherency unit if one of the records indicates that no active device in the node has the ownership responsibility for the coherency unit and the coherency unit is in the modified global access state in the node.

5

3. The node of claim 2, wherein the one of the records corresponds to a write back address packet sent to initiate a write back transaction for the coherency unit, wherein the coherency unit does not map to any memory subsystem included in the node.

10

4. The node of claim 3, wherein the address network is configured to convey the write back address packet in broadcast mode, and wherein one of the plurality of active devices that sends the write back address packet is configured to lose the ownership responsibility for the coherency unit in response to receiving the write back address packet.

15

5. The node of claim 3, wherein in response to the one of the records, the interface is configured to send a pull request to one of the plurality of active devices that initiated the write back transaction;

20

wherein in response to the pull request, the one of the plurality of active devices is configured to send the data corresponding to the coherency unit to the interface.

25

6. The node of claim 5, wherein the one of the plurality of active devices is configured to transition an access right to the coherency unit upon sending the data.

30

7. The node of claim 5, wherein the active device is configured to send the address packet to initiate a read-to-own transaction, and wherein the interface is configured to send the data corresponding to the coherency unit to the active device in response to accessing one of the plurality of records corresponding to the address packet and in

response to receiving the data from the one of the plurality of active devices that initiated the write back transaction.

8. The node of claim 7, wherein the active device is configured to gain an ownership
5 responsibility for the coherency unit in response to receiving the address packet, wherein the interface is configured to send a proxy address packet on the address network in response to receiving the request from the other node, wherein the active device is configured to store a promise corresponding to the proxy address packet in a promise array included in the active device in response to receiving the proxy address packet
10 while having the ownership responsibility for the coherency unit.

9. The node of claim 8, wherein in response to the promise, the active device is configured to send data corresponding to the coherency unit to the interface, and wherein the interface is configured to supply the data to the other node in response to the request
15 upon receiving the data from the active device.

10. The node of claim 2, wherein the one of the records corresponds to a write stream address packet sent to initiate a write stream transaction for the coherency unit.

20 11. A system, comprising:

a node comprising a plurality of active devices, an interface to an inter-node network, and an address network configured to convey address packets between the interface and the plurality of active devices;

25

an additional node coupled to the node by the inter-node network;

wherein an active device of the plurality of active devices is configured to send an address packet on the address network to initiate a transaction to gain an
30 access right to a coherency unit;

wherein in response to the address packet, the interface is configured to send data corresponding to the coherency unit to the active device if no other active device in the node has an ownership responsibility for the coherency unit and the coherency unit is in a modified global access state in the node.

12. The system of claim 11, wherein the interface includes storage for a plurality of records, wherein each of the plurality of records corresponds to a respective one of a plurality of address packets conveyed on the address network;

wherein the interface is configured to search the storage for records specifying the coherency unit in response to the interface receiving a request for access to the coherency unit from the additional node via the inter-node network;

wherein the interface is configured to send the data corresponding to the coherency unit if one of the records indicates that no active device in the node has the ownership responsibility for the coherency unit and the coherency unit is in the modified global access state in the node.

13. The system of claim 12, wherein the one of the records corresponds to a write back address packet sent to initiate a write back transaction for the coherency unit, wherein the coherency unit does not map to any memory subsystem included in the node.

14. The system of claim 13, wherein the address network is configured to convey the write back address packet in broadcast mode, and wherein one of the plurality of active devices that sends the write back address packet is configured to lose the ownership responsibility for the coherency unit in response to receiving the write back address packet.

15. The system of claim 13, wherein in response to the one of the records, the interface is configured to send a pull request to one of the plurality of active devices that initiated the write back transaction;

5 wherein in response to the pull request, the one of the plurality of active devices is configured to send the data corresponding to the coherency unit to the interface.

16. The system of claim 15, wherein the one of the plurality of active devices is
10 configured to transition an access right to the coherency unit upon sending the data.

17. The system of claim 15, wherein the active device is configured to send the address packet to initiate a read-to-own transaction, and wherein the interface is configured to send the data corresponding to the coherency to the active device in
15 response to accessing one of the plurality of records corresponding to the address packet and in response to receiving the data from the one of the plurality of active devices that initiated the write back transaction.

18. The system of claim 17, wherein the active device is configured to gain an
20 ownership responsibility for the coherency unit in response to receiving the address packet, wherein the interface is configured to send a proxy address packet on the address network in response to receiving the request from the additional node, wherein the active device is configured to store a promise corresponding to the proxy address packet in a promise array included in the active device in response to receiving the proxy address
25 packet while having the ownership responsibility for the coherency unit.

19. The system of claim 18, wherein in response to the promise, the active device is configured to send data corresponding to the coherency unit to the interface, and wherein the interface is configured to supply the data to the additional node in response to the
30 request upon receiving the data from the active device.

20. The system of claim 12, wherein the one of the records corresponds to a write back address packet sent to initiate a write stream transaction for the coherency unit.

5 21. A method of operating a multi-node system comprising a node coupled to an additional node by an inter-node network, the method comprising:

an active device of a plurality of active devices included in the node sending an address packet on the address network to initiate a transaction to gain an
10 access right to a coherency unit;

in response to the address packet, an interface to the inter-node network included in the node sending data corresponding to the coherency unit to the active device if no other active device in the node has an ownership
15 responsibility for the coherency unit and the coherency unit is in a modified global access state in the node.

22. The method of claim 21, further comprising:

20 the interface storing a plurality of records, wherein each of the plurality of records corresponds to a respective one of a plurality of address packets conveyed on the address network;

in response to the interface receiving a request for access to the coherency unit
25 from the additional node via the inter-node network, the interface searching the plurality of records for records specifying the coherency unit;

the interface sending the data corresponding to the coherency unit if one of the records specifying the coherency unit indicates that no active device in the

node has the ownership responsibility for the coherency unit and the coherency unit is in the modified global access state in the node.

23. The method of claim 22, wherein the one of the records corresponds to a write
5 back address packet sent to initiate a write back transaction for the coherency unit,
wherein the coherency unit does not map to any memory subsystem included in the node.

24. The method of claim 23, further comprising:
10 the address network conveying the write back address packet in broadcast mode;
and
one of the plurality of active devices that sent the write back address packet losing
the ownership responsibility for the coherency unit in response to
15 receiving the write back address packet.

25. The method of claim 23, further comprising:
in response to the one of the records, the interface sending a pull request to one of
20 the plurality of active devices that initiated the write back transaction;
in response to the pull request, the one of the plurality of active devices sending
the data corresponding to the coherency unit to the interface.

25 26. The method of claim 25, further comprising the one of the plurality of active
devices transitioning an access right to the coherency unit upon sending the data.

27. The method of claim 25, wherein said the active device sending the address
packet comprises the active device sending the address packet to initiate a read-to-own
30 transaction;

wherein said the interface sending the data corresponding to the coherency unit to the active device occurs in response to the interface accessing one of the plurality of records corresponding to the address packet and receiving the data from the one of the plurality of active devices that initiated the write back transaction.

28. The method of claim 27, further comprising:

the active device gaining an ownership responsibility for the coherency unit in response to receiving the address packet;

the interface sending a proxy address packet on the address network in response to receiving the request from the additional node; and

the active device storing a promise corresponding to the proxy address packet in a promise array included in the active device in response to receiving the proxy address packet while having the ownership responsibility for the coherency unit.

29. The method of claim 28, further comprising:

in response to the promise, the active device sending data corresponding to the coherency unit to the interface; and

the interface supplying the data to the additional node in response to the request upon receiving the data from the active device.

30. The method of claim 22, wherein the one of the records corresponds to a write back address packet sent to initiate a write stream transaction for the coherency unit.